

# Fine tuning embeddings for nuclear power

Rob Whelan  
Gridway AI



June 4, 2025

AIMUG, Austin, TX

---

<https://github.com/gridwayai/gridwayai-sdk>

<https://huggingface.co/gridwayai/nuclear-licensing-embeddings-768>

# Content

1 Nuclear has lots of jargon >

2 Better embeddings, better search >

3 Coolant is not just for your fridge >

4 How to process training data >

5 Code >

6 Huggingface & deployment >

LWRs classified as advanced reactors under NEIMA are more likely to resemble existing NRC-licensed LWRs than non-LWRs, particularly in terms of design, operating experience, and the applicability of current regulations and guidance. NRC staff are therefore leveraging much of the existing regulatory guidance for LWRs as a starting point for reviews of LWRs that could be classified under NEIMA as advanced reactors. As a result, while both LWR and non-LWR designs can be defined as advanced reactors under NEIMA, the NRC's definition of advanced reactors generally applies only to non-LWR designs unless otherwise specified. The implications and implementation of this regulatory distinction are discussed further below.

example

Acronyms galore

Some represent an entire universe of meaning - Light Water Reactors (LWR) means something TOTALLY DIFFERENT than "GEN4+" reactors.

# Nuclear has lots of jargon

# Better embeddings → better search


Embeddings, or vector representations of words and phrases, are arrays of floats, like [0.133, -1.533, 2.122, 0.001,...]

You do math to find out how close vectors are to each other – cosine similarity, or Euclidean distance


If your embeddings understand the jargon and phrasing of your material, you're going to get better search results.

The way we speak in text messages is very different from academic language, for example

# Before fine-tuning



# After



# The notebook, on AWS

ml.g6.16xlarge

You need a GPU and plenty of memory, for pytorch and to store the base embeddings model

80% train, 20% validation

10,000 examples in  
embedding\_data\_hard\_negs\_4.jsonl -  
generated by gpt-4o-mini using a few dozen  
regulatory source texts

<https://github.com/gridwayai/gridwayai-sdk>

'BAAI/bge-base-en-v1.5' (768)

Base embeddings model, widely used as an English embeddings model, probably slightly less performant than `text-embedding-3-small`, with dimension 1536

MultipleNegativesRankingLoss

Pass in positive and negative pairs; do some "hard negatives" which are "hard" to differentiate, but really are different

<https://huggingface.co/gridwayai/nuclear-licensing-embeddings-768>

# Use the model

```
from sentence_transformers import SentenceTransformer, util

model = SentenceTransformer("gridwayai/nuclear-licensing-embeddings-768")

sentences = [
    'What is the purpose of the Rapid Borate Stop Valve in Reactor Control?',
    'Locates and discusses opening 1CV175, Rapid Borate Stop Valve by disengaging clutch and rotating handwheel (counterclockwise).',
    'CLOSE the Air Supply Isolation Valve, 12CV160 A/S, AIR SUPPLY FOR 12CV160.',
]
embeddings = model.encode(sentences)

→ returns a list of vector arrays
```